

Amendments to the Specification:

Please replace with the Abstract of the Disclosure with the following Abstract of the Disclosure.
A copy of the Abstract of the Disclosure is annexed hereto on a separate sheet as required.

~~The present invention provides an~~ An audio recording/playback tool ~~that is integrated with an~~
information viewer that simplifies recording and playback of audio annotations. ~~The invention~~
~~also provides alternative~~ Alternative techniques are provided to retrieve, categorize and sort the
audio annotations including the ability to associate audio annotations with either pages of a
document or specific points inside a page. Further, ~~the invention synchronizes audio playback~~
and document navigation actions can be synchronized. ~~The invention supports the storage of the~~
~~audio~~ Audio annotations can be stored in a variety of formats including ~~the annotations stored as~~
discrete clips labeled with properties and ~~stored in an external database that permits, among other~~
things, exchanging of annotations between users.

Please replace the paragraph beginning at page 5, line 17, with the following amended paragraph:

Figure 3 is a representation of a screen ~~of~~ having a simplified audio annotation interface
according to embodiments of the invention.

Please replace the paragraph beginning at page 5, line 19, with the following amended paragraph:

Figure 4 is a representation of a screen ~~of~~ having an advanced audio annotation interface
according to embodiments of the invention.

Please replace the paragraph beginning at page 11, line 17, with the following amended paragraph:

Audio annotations are combinations of one or more audio clips. As a user speaks, the system
recording the user's voice stores received information as audio clips. The audio clips are
separated from each other based a variety of events including: 1) momentary pauses in the user's

speech, 2) user actions on the device, such as navigating between pages or documents, and 3) timeouts that set the maximum duration of a clip if neither 1 nor 2 occurs first. The user may be unaware of the fact that annotations are stored as sets of clips. On playback, the system assembles the clips into audio annotations. By the system forming annotations from stored audio clips, the system is able to make finer resolutions between spoken comments (for example when a user continues to speak across numerous pages). These finer resolutions are helpful in interpolating when annotations are to be separated for various purposes including purposes of editing (insert/delete) or playback indexing. By means of example, the system may record a user's voice as a first file, and then parses the file to extract the audio clips. As is appreciated by one of ordinary skill in the art ~~that~~ the parsing may occur in real time, may be performed while no speech is occurring (during processor down time), or may be uploaded for processing at a later time.

Please replace the paragraph beginning at page 13, line 14, with the following amended paragraph:

Properties are associated with audio clips when created and/or when stored as described above. Properties help a user retrieve audio clips as audio annotations. The audio clips may be stored in a database to facilitate dynamically accessing the audio clips based on user-defined queries. This ability to retrieve the audio information based on user input is a separation from the linear nature of recording ~~that most users' users~~ expect. Here, the storage of the audio information includes properties that permit the audio information to be associated with ~~the-visual~~ information so that one may be ~~displayed-generated~~ in synchronism with the other.

Please replace the paragraph beginning at page 15, line 6, with the following amended paragraph:

Figure 2 is schematic representation of a set of audio clips 202. The set of audio clips 202 is typically formed of multiple individual audio clips that have been separately recorded. Any number of audio clips may be associated with any page of textual information. In addition, the audio clips may be recorded at a variety of different times. The electronic information (shown

here as pages) in Figure 2 are provided as pages in an electronic book. Once inserted, the audio clips add richness to textual electronic information. On playback, the set of audio clips may be combined into a single audio stream and ~~is~~can be derived by query from a database. It is appreciated that any type of electronic information, for example video, may be displayed on any device supporting electronic reading. In the example of annotating video information, adding audio annotations to a video presentation permits a user to comment on displayed video information.

Please replace the paragraph beginning at page 15, line 21, with the following amended paragraph:

In the present example, individual audio clips 202a through 202n comprise audio clip set 202. As shown in the example of Figure 2, the audio clips may be stored as individual audio notes or portions that may be arranged into audio annotations based on user preference. For example, Figure 2 shows individual audio clips being associated with pages of a first book 204 and pages of a second book 206. More specifically, two individual audio clips 202a and 202b are associated with page 10 of the first book 204; one clip 202c is associated with page 11 of first book 204, etc. Other individual audio clips are associated with second book 206. In the example, page 56 of book 206 has associated audio clips 202h, 202i and 202j. In one embodiment, the process of selecting individual audio clips 202a through 202n to the set of audio clips 202 is transparent to the user. For example, a user may request all audio clips associated with Book 1 be sorted in page order. The resulting audio stream would include audio clips 202a-202g. In another embodiment, the user may request all audio annotations for Books 1 and 2 in order of recording time recorded before a given date. The resulting audio stream may include, for example, the following clips in order: from Book 1, 202a, 202d, 202b, 202c, 202e, then flipping to Book 2, clips 202h, 202k, 202i, 202l, then back to Book 1 for clips 202g and 202f. Here, clips 202j, 202m, and 202n may have been recorded after the given date. In a third example, a user may request all audio clips be arranged in relation to the author or content of the comment including "all audio clips by Mr. Jones" or "all audio clips relating to astronomy". In regards to the content, the system may include a property in the audio clips that defines the content. This may

be accomplished as well as by the title of the audio clip or by the title of the viewed document as stored with the audio clip when the audio clip was made. In short, the order of the audio clips in the audio stream is dependent on how a user queries a database (where the database storage structure is used). Further, predefined queries may also exist that permit a user with canned playback orders, thus minimizing the number of separate inputs a user has to make to start playback. Examples of the canned queries include "all annotations of currently viewed document, ascending in creation time order", "all annotations of all documents, descending in creation time order", etc. Other combinations and permutations for stored queries are possible and considered within the scope of the invention.

Please replace the paragraph beginning at page 17, line 22, with the following amended paragraph:

A user may concurrently access a number of tapes while reading a document. For instance, a user may have a first tape for notes on the content of a book, have a second tape for notes of additional books the user would wish to read, have a third tape for adding editorial comments for another user, have a ~~fifth~~ fourth tape for recording audio annotations taken in conjunction with a presentation, and have a ~~third~~ fifth tape (unrelated to the first two tapes) for recording of notes of items to pick up at the grocery store after getting home. In this regard, selecting a tape then recording generates audio clips with properties including the user's current focus, including, at least in part, the name or other identifier of the selected tape.

Please replace the paragraph beginning at page 18, line 7, with the following amended paragraph:

As applied to Figure 3, display portion 310 indicates the identity of the tape currently receiving/playing back audio annotations. It is appreciated that the identity of the tape is definable by the user. The ability to name tapes makes ~~for~~ later identification easier. The names may relate to previous queries. For example, a user may have a tape named "History Class Notes" where the database query was "all annotations where subject is 'history class'". In another embodiment, the system also provides intelligent naming of audio clips to match that of

the tape currently being recorded or played back. For example, when playing back a tape "History Class Notes", a user may create a new audio annotation to comment on a previous audio note. Here, the system determines the name of the current tape "History Class Notes" and assigns properties to the new audio clip to make it part of the History Class Notes tape. In the example of the audio notes being stored in a database, the new audio clip would have the property "subject=history class" so as to be part of the History Class Notes tape (or, more precisely, the virtual tape or audio stream) as described above. The property may be represented in a number of forms including XML and other mark up languages or by a predefined coding system and the like.

Please replace the paragraph beginning at page 22, line 18, with the following amended paragraph:

Where desired, play and fast forward or rewind may be engaged simultaneously. This simulates the operation of a physical tape. Here, the system may use a compression algorithm to play back an excerpted version of the audio version of the audio stream as the tape winds. Alternatively, the audio annotation may be rendered in a high pitch, providing the modulations of the recorded voice, but at a fast rate. Thus, audio cues are provided about where the tape is positioned. To repeat what was just listened to or recorded, a button may be pressed for playback. Playback or recording resumes after the repeated interval. All tapes (including master tapes, document tapes, and any other predefined or executed queries) may be scanned, played, or have material appended thereto. Recording at the end of the tape appends the new clips to the tape.

Please replace the paragraph beginning at page 23, line 16, with the following amended paragraph:

A settings sheet (not shown for simplicity) allows the user to preset various features of the device, such as to ~~inactivate~~ deactivate the locking behavior of the fast forward and rewind buttons relating to a user's preferences. Similar settings may include determining the speed of fast forwarding and rewinding.

Please replace the paragraph beginning at page 23, line 20, with the following amended paragraph:

In one aspect of the present invention, the controls for the system are normally not visible until implemented by a toolbar that is, by default, generally hosted in a command shortcut margin and initially closed. In this implementation, a toolbar tab is found in the shortcut margin, similar to a bookmark tab. Activating the tab opens the interface portion 403 (or 303) into the margin. In one implementation, the toolbar slides out ~~from~~ from the margin edge. Activating the tab again retracts it, leaving only the tab. For convenience, where desired, the toolbar may be deleted or moved to a different desired location. Where the toolbar tab has been deleted, it may be recovered by obtaining another copy of the toolbar as is known in the art.

Please replace the paragraph beginning at page 27, line 6, with the following amended paragraph:

Figure 6 shows an example of a user note that may contain an audio annotation as reflected by icon 415 of Figure 4. In addition to being able to associate audio clips with pages or items in a viewed document, the system permits audio information to be associated with text notes or other displayed ~~item~~ items or information. For example, a document author may create a document with a link between a word, a graphic image, or an icon and an audio annotation. So, by tapping the item (word, graphic image or icon), the link is activated and the system plays the related audio annotation. Figure 6 shows a text note 601 on page 600 with an audio annotation 602 associated with the text note 601. The audio annotation represented by icon 602 may start to play automatically after a user accesses note 601 or may wait for a user to tap on it prior to playing.

Please replace the paragraph beginning at page 29, line 21, with the following amended paragraph:

One distinction between the second implementation versus the first one ~~lined~~ is that the second implementation is simpler and has more features. That is, rather than have one mechanism for associating audio clips with ranges of document positions (for page-level audio) and another one for ~~associate~~ associating audio clips with embedded links, the system uses page-level audio only

and ~~take~~takes advantage of the another existing feature (embedded notes) to provide the functionality of a link to audio. That is, from the user's point of view, the behavior is the same—tap an icon and audio plays. But the second mechanism is simpler (one mechanism instead of two) and more powerful (because one may always add ink/text to the audio note, or go back to an ink/text note and add audio, and thus have notes that contain both media).

Please replace the paragraph beginning at page 32, line 21, with the following amended paragraph:

Automatic playback (also referred to as single touch playback) enables a mode of reading a document and reviewing recorded notes where a user simply points at notes to hear their associated audio content. In other words, imagine a person as they read along, simply tapping this note and then that note to hear its content. The importance of this feature is that it makes the process of reviewing the audio content of notes very transparent so that it does not interfere with or slow down the process of reading the document. It's also significant that there are different cases of note playback here. One is tapping on an embedded note, in which case that note's content is played back. Another is that of tapping on an overlaid note, such as some handwriting in the margin of the document, or a stretch of highlighted text. What happens in this case is that the audio that is played back is the audio that was recorded in association with that page of the document at the same time as when that note was ~~been~~ entered onto the page. For example, imagine a lecture presentation with slides, and one reviews the slides later with notes one wrote on the slides. By tapping on any of the notes, one is able to hear what the lecturer was saying at the point in time when one was writing the note. As with the embedded note case, auto playback makes it very simple to read through the set of slides and retrieve the relevant audio context associated with each of the notes one scribbled.

Please replace the paragraph beginning at page 39, line 14, with the following amended paragraph:

Figure 13 describes a process for adding information to a document. First, in step 1301, the system receives a user request to add information. The user may want to add a written annotation

(ink, highlights, underlining and the like) or add audio. This request may come in the form of speaking, tapping on a screen, writing on a screen, tapping a link, or the like. The system creates a link object in step 1302 to associate the information to be added with the document. In step 1303, the system adds information relating to the source document to the link object as the source anchor. The source anchor may ~~including~~ include the name of the document, for example, "source document name = host doc 1". The source anchor may include other properties as described above.

Please replace the paragraph beginning at page 41, line 4, with the following amended paragraph:

In reference to Figure 13, it is noted that, if there are embedded notes on a page, one may tap on them to play back their contained audio (if any) or one may create and speak into new embedded notes. Here, the system simply changes what set of properties it is using to retrieve or store audio clips. As a result, one is free to create an embedded note that will contain both audio and text, or that will start out with text only at first and adds audio ~~to~~-later, or that starts out ~~as~~-with audio only and ~~you~~-later adds text-~~to~~.